arXiv:1512.09349v1 [cs.DS] 31 Dec

Faster Maximium Priority Matchings in Bipartite Graphs

Jonathan Turner

WUCSE-2015-08

Abstract

A maximum priority matching is a matching in an undirected graph that maximizes a *priority score* defined with respect to given vertex priorities. An earlier paper showed how to find maximum priority matchings in unweighted graphs. This paper describes an algorithm for bipartite graphs that is faster when the number of distinct priority classes is limited. For graphs with k distinct priority classes it runs in $O(kmn^{1/2})$ time, where n is the number of vertices in the graph and m is the number of edges.

The maximum priority matching problem was introduced in [9]. In this problem, each vertex has an integer-valued *priority* and the objective is to find a matching that maximizes a *priority score* defined with respect to these values. The priority score is defined as the n-ary number in which the i-th most-significant digit is the number of matched vertices with priority i. The earlier paper described an algorithm for finding maximum priority matchings in $O(mn)$ time that is based on Edmonds' algorithm for the maximum size matching problem [2, 3, 6].

In a private communication, Tarjan observed that the 2-priority case for bipartite graphs could be solved in $O(mn^{1/2})$ time using a recent algorithm for weighted matchings in bipartite graphs [1, 4, 7]. One assigns a weight of 0, 1 or 2 to each edge based on the number of high priority vertices it is incident to. A maximum weight matching of this graph matches the largest possible number of high priority vertices. It can then be extended to a maximum priority matching using an algorithm by Hopcroft and Karp [5]. Tarjan speculated that this approach might be extended to handle multiple priority classes. This paper takes a different approach, yielding a simpler algorithm that handles the k-priority case in $O(kmn^{1/2})$ time, making it faster than the earlier algorithm when k grows more slowly than $n^{1/2}$.

We assume that the reader is familiar with [9], which describes a generalization of the augmenting path algorithm of Edmonds. In this algorithm there are two types of augmenting paths. Odd-length augmenting paths increase the number of edges in the current matching and may also increase its priority score, while even-length augmenting paths increase the priority score without changing the size of the matching.

Here, we take a slightly different approach. Let $G = (V, E)$ be an undirected bipartite graph with vertex priorities $p(u)$ and let V_1 and V_2 define the bipartition of G (that is, all edges join vertices in V_1 to vertices in V_2). The algorithm starts by finding a maximum size matching, then sets an integer variable i to the index of the first non-empty priority class and repeats the following step until all non-empty priority classes have been processed.

> While there are even-length i-augmenting paths with an unmatched vertex in V_1, find such a path and reverse the matching status of its edges. While there are even-length i-augmenting paths with an un-

matched vertex in V_2, find such a path and reverse the matching status of its edges. Advance i to the index of the next non-empty priority class.

The augmenting path searches in each step can be implemented by solving a pair of maximum flow problems on *unit graphs*. These can be solved in $O(mn^{1/2})$ time using Dinic's algorithm [8]. The original maximum size matching can also be found in $O(mn^{1/2})$ time using the Hopcroft-Karp algorithm, yielding an overall time bound of $O(kmn^{1/2})$ for k priority classes.

Let M_1 be the matching at the start of step i. We construct an instance of the maximum flow problem $X_1 = (W_1, F_1)$ as follows.

$$\begin{aligned} W_1 &= V \cup \{s, t\} \\ F_1 &= \{(u, v) \mid u \in V_1,\ v \in V_2 \text{ and } \{u, v\} \notin M_1\} \\ &\cup \{(v, u) \mid u \in V_1,\ v \in V_2 \text{ and } \{u, v\} \in M_1\} \\ &\cup \{(s, u) \mid u \in V_1 \text{ is unmatched and } p(u) = i\} \\ &\cup \{(u, t) \mid u \in V_1 \text{ is matched and } p(u) > i\} \end{aligned}$$

All edges are assigned a capacity of 1. Observe that X_1 is a unit graph, since each vertex has at most one incoming or one outgoing edge. A maximum flow f_1 on X_1 can be decomposed into a set of augmenting paths in X_1. Each of these paths corresponds directly to an i-augmenting path in G and consequently, f_1 defines a modified matching M_2.

$$\begin{aligned} M_2 &= \{\{u, v\} \mid u \in V_1,\ v \in V_2,\ (u, v) \in F_1 \text{ and } f_1(u, v) = 1\} \\ &\cup \{\{u, v\} \mid u \in V_1,\ v \in V_2,\ (v, u) \in F_1 \text{ and } f_1(v, u) = 0\} \end{aligned}$$

Note that every vertex with priority $\leq i$ that is matched in M_1 is also matched in M_2.

Next, the algorithm constructs a second flow problem $X_2 = (W_2, F_2)$.

$$\begin{aligned} W_2 &= V \cup \{s, t\} \\ F_2 &= \{(u, v) \mid u \in V_1,\ v \in V_2 \text{ and } \{u, v\} \notin M_2\} \\ &\cup \{(v, u) \mid u \in V_1,\ v \in V_2 \text{ and } \{u, v\} \in M_2\} \\ &\cup \{(s, v) \mid v \in V_2 \text{ is matched and } p(u) > i\} \\ &\cup \{(v, t) \mid v \in V_2 \text{ is unmatched and } p(u) = i\} \end{aligned}$$

All edges are again assigned a capacity of 1. A maximum flow f_2 on X_2 defines a modified matching M_3.

$$\begin{aligned} M_3 &= \{\{u,v\} \mid u \in V_1,\, v \in V_2,\, (u,v) \in F_2 \text{ and } f_2(u,v) = 1\} \\ &\quad \cup \{\{u,v\} \mid u \in V_1,\, v \in V_2,\, (v,u) \in F_2 \text{ and } f_2(v,u) = 0\} \end{aligned}$$

Observe that every vertex with priority $\leq i$ that is matched in M_1 is also matched in M_3.

References

[1] Duan, Ran, Seth Pettie, and Hsin-Hao Su. "Scaling algorithms for approximate and exact maximum weight matching," *Computing Research Repository* (CoRR), abs/1112.0790, 2011.

[2] Edmonds, Jack. "Paths, trees and flowers," *Canadian Journal of Mathematics*, 1965, pp. 449–467.

[3] Gabow, Harold N. "An efficient implementation of Edmonds' algorithm for maximum matching on graphs," *Journal of the Association for Computing Machinery*, 1976, pp. 221–234.

[4] Goldberg, Andrew V., Haim Kaplan, Sagi Hed, and Robert E. Tarjan. "Minimum cost flows in graphs with unit capacities," *Symposium on Theoretical Aspects of Computer Science* (STACS), 2015.

[5] Hopcroft, John E. and Richard M. Karp. "An $O(n^{5/2}$ algorithm for maximum matching in bipartite graphs," *SIAM Journal on Computing*, 1973, pp 225–231.

[6] Micali, Silvio. and V. V. Vazirani. "An $O(\sqrt{|V|} \cdot |E|)$ algorithm for finding maximum matchings in general graphs," *IEEE Symposium on the Foundations of Computer Science (FOCS)*, 1980, pp. 17-27.

[7] Ramshaw, L., and Robert E. Tarjan. "A weight-scaling algorithm for min-cost imperfect matchings in bipartite graphs". In *IEEE Symposium on the Foundations of Computer Science*, pages 581590, 2012.

[8] Tarjan, Robert E. *Data structures and network algorithms*. Society for Industrial and Applied Mathematics, 1983.

[9] Turner, Jonathan S. "Maximum priority matchings," Washington University Computer Science and Engineering Department technical report, WUCS-2015-06, 2015.